Eco514—Game Theory
Lecture 2: (Iterated) Best Response Operators

Marciano Siniscalchi

September 21, 1999

Introduction

This lecture continues the analysis of normal-form games. We analyze general, non-zerosum games, emphasizing the informal "equation":

Rational Behavior + Assumptions about Beliefs = Solution Concepts

Before we tackle the new material, let us review what we have learned about zerosum games in light of this "equation". *Rational behavior* in the context of normal-form games (zerosum or otherwise) means that players choose strategies which maximize their expected payoff, given their beliefs about their opponents' choices. Today we shall examine a few technical aspects related to the existence and characteristics of the *best reply correspondence*, but the basics should be familiar to you from decision theory.

The minmax theory makes the following *assumption about beliefs*: each player expects her opponent to (1) understand that she will best-respond to his strategy, and (2) play so as to minimize her expected payoff. That is, each player reasons *as if* she were to *observe* the (possibly random) choice of her opponent and best-respond to it, and *as if* her opponent, *anticipating this*, chose a strategy in order to make her payoff as low as possible.

Formally, the problem

$$\min_{\alpha_j^i \in \Delta(A_j)} \max_{\alpha_i^i \in \Delta(A_i)} u_i(\alpha_i^i, \alpha_j^i)$$

defines the belief α_j^i of Player i about Player j.

Best Replies

Let us begin with a formal definition of best replies in the case of finite games.

VISIT…

Definition 1 *Fix a finite normal-form game $G = (N, (A_i, u_i)_{i \in N})$, a player $i \in I$ and a probability measure $\alpha_{-i} \in \Delta(A_{-i})$. An action $a_i \in A_i$ is a* best reply *to the belief α_{-i} iff*

$$\sum_{A_{-i}} \alpha_{-i}(a_{-i}) u_i(a_i, a_{-i}) \geq \sum_{A_{-i}} \alpha_{-i}(a_{-i}) u_i(a'_i, a_{-i}) \quad \forall a'_i \in A_i.$$

In finite games, expected payoffs are always well-defined. Not so for games with infinite action spaces: even assuming a nice measure-theoretic structure, it is not hard to construct (mildly pathological) examples in which payoff functions are not integrable.

Thus, a better definition for infinite games is the following. Assume that, for every player $i \in N$, the set of actions $(A_i, \mathcal{A}_i)$ is a measure space. Endow the product sets A and A_{-i} with the natural product sigma-algebras, denoted $\mathcal{A}$ and $\mathcal{A}_{-i}$. Finally, assume that u_i is $\mathcal{A}$-measurable. The sigma-algebras on the set of actions will be indicated explicitly for infinite games.

Definition 2 *Fix a normal-form game $G = (N, (A_i, \mathcal{A}_i, u_i)_{i \in N})$, a player $i \in N$ and a belief $\alpha_{-i} \in \Delta(A_{-i}, \mathcal{A}_{-i})$. An action $a_i \in A_i$ is a* best reply *to the belief α_{-i} iff the Lebesgue integral $\int_{A_{-i}} u_i(a_i, a_{-i}) \alpha_{-i}(da_i)$ exists, and moreover*

$$\int_{A_{-i}} u_i(a_i, a_{-i}) \alpha_{-i}(da_i) \geq \int_{A_{-i}} u_i(a'_i, a_{-i}) \alpha_{-i}(da_i) \quad \forall a'_i \in A_i$$

whenever the right-hand side integral exists.

I emphasize that integrals are taken (here and in the following) in the Lebesgue sense.

Even if expected payoffs are well-defined, the very structure of the game might lead to the non-existence of best replies. An obvious example arises in the Bertrand price competition game.

Example. Two firms produce and sell the same good, facing a market of size Q, and zero fixed and marginal cost; finally, the two firms compete by posting unit prices p_i, $i = 1, 2$. Assume that consumers have unit demand and buy from the cheapest producer; if $p_1 = p_2$, consumers buy from each firm with probability $\frac{1}{2}$.

Suppose that Firm 1 expects Firm 2 to post a price equal to $p_2 = 2$; this corresponds to a degenerate belief, concentrated on 2. Then any price $p_1 > 2$ yields 0, any price $p_1 < 2$ yields Qp_1, and $p_1 = p_2 = 2$ yields $\frac{2Q}{2} = Q$. Clearly no price p_1 satisfies our definition of a best reply.

The conditions for the existence of best replies are related to those which guarantee the existence of an optimum. We need to specify a topology $\mathcal{T}_i$ on every action space A_i;

furthermore, we assume that the relevant sigma-algebra in the definition of a best reply is the corresponding Borel sigma-algebra $\mathcal{B}(\mathcal{T}_i)$ (hence, $u_i(\cdot)$ is Borel-measurable if it is continuous). Finally, we assume product sets are endowed with the product topology and the induced Borel product sigma-algebra.

Proposition 0.1 *Consider a game $G = (N, (A_i, \mathcal{T}_i, u_i)_{i \in N})$ and fix a player $i \in N$. If $u_i(\cdot)$ is bounded and continuous, and A_i is a compact metrizable space, then for every measure $\alpha_{-i} \in \Delta(A_{-i}, \mathcal{B}(\mathcal{T}_i))$ there exists a best reply $a_i \in A_i$.*

Proof: We need to show that the map $\varphi : A_i \mapsto \int_{A_{-i}} u_i(a_i, a_{-i}) \alpha_{-i}(da_{-i})$ is continuous. Since A_i is a metric space, sequences characterize continuity. Thus, fix a sequence $a_i^k \to a_i$ in A_i. By assumption, $u_i(a_i^k, a_{-i}) \to u_i(a_i, a_{-i})$ for every $a_{-i} \in A_{-i}$, and $u_i(\cdot)$ is bounded. Thus, by Lebesgue's Dominated Convergence theorem, $\varphi(a_i^k) \to \varphi(a_i)$. ■

Observe that, in the above proof, it is essential that expectations be defined in the Lebesgue sense: otherwise, stronger conditions are required on the payoff functions and action spaces.

There are two reasons why we are interested in infinite games. First, for Nash equilibria to exist one must enlarge the players' action spaces to include all mixtures of strategies; thus, effectively, one is looking at an infinite game, in which every strategy set is a simplex, hence a compact subset of Euclidean space, and every payoff function is continuous (because the expected payoff from a mixed strategy profile is linear in the probabilities). Thus, the preceding result applies.

Second, several interesting economic models lend themselves to a formulation involving infinite strategy spaces. Thus, getting some exposure to the issues that arise in such settings can be useful.

The Best Reply Correspondence

Recall that a *correspondence* is a set-valued function. Since, in general, Player i may have more than one best reply to a belief $\alpha_{-i} \in \Delta(A_{-i}, \mathcal{A}_{-i})$, we use correspondences to describe the mapping from beliefs to best replies.

Definition 3 *Fix a game $G = (N, (A_i, \mathcal{A}_i)_{i \in N})$ and a player $i \in N$. The* best-reply correspondence *for player i, $r_i : \Delta(A_{-i}, \mathcal{A}_{-i}) \Rightarrow A_i$, is defined by*

$$\forall \alpha_{-i} \in \Delta(A_{-i}, \mathcal{A}_{-i}), \quad a_i \in r_i(\alpha_{-i}) \Leftrightarrow a_i \text{ is a best reply given } \alpha_{-i}$$

The preceding proposition then gives conditions under which the best-reply correspondence is nonempty.

Correlated Rationalizability and Iterated Dominance

Having disposed of all required technicalities, let us go back to our "informal equation"

Rational Behavior + Assumptions about Beliefs = Solution Concepts.

Fix a game $G = (N, (A_i, \mathcal{T}_i, u_i)_{i \in N})$ and assume that each $(A_i, \mathcal{T}_i)$ is a compact metrizable space (you can also assume that each A_i is finite, as far as the substantive arguments are concerned) and that payoff functions are continuous.

We begin by asking: what is a *plausible outcome of the game*, if we are willing to assume that players are (Bayesian) rational? The answer should be clear: we can only expect an action profile $(a_i)_{i \in N}$ to be chosen iff, for every player $i \in N$, $a_i \in r_i(\alpha_{-i})$ for some $\alpha_{-i} \in \Delta(A_{-i}, \mathcal{B}(\mathcal{T}_i))$.

That is: if we think that players are "good Bayesians," we should expect them to formulate a belief about their opponents' choice, and play a best reply given that belief. Equivalently, we should expect a rational player to choose an action only if it that action is *justified* by some belief.

This exhausts *all* the implications of Bayesian rationality: in particular, *we do not postulate any relationship between the beliefs justifying the actions in any given profile.*

In very simple games, Bayesian rationality is sufficient to isolate a unique solution: for instance, this is true in the Prisoner's Dilemma. However, in Matching Pennies, the Row player (who wants to match) may choose H because she expects the Column player to also choose H, whereas the Column player (who wants to avoid matches) actually chooses T because he (correctly, it turns out) expects the Row player to choose H. Indeed, it is easy to see that any action profile is consistent with Bayesian rationality in Matching Pennies.

However, one might reason as follows:

> If I, the theorist, am able to rule out certain actions because they are never best-replies, perhaps the players will be able to do so, too. But, if so, it seems *reasonable* to assume that their beliefs should assign zero probability to the collection of eliminated actions. *This entails a further restriction on the collection of actions that they might choose.*

That is, in addition to assuming that players are Bayesian rational, we may be willing to assume that *they believe that their own opponents are also Bayesian rational.* This is precisely the type of assumption about beliefs referred to in our "informal equation." It is also a very powerful idea:

Example: Consider a two-firm version of the Cournot quantity competition game, with action spaces $A_i = [0, 1]$ and payoffs $u_i(a_i, a_{-i}) = (2 - a_i - a_{-i})a_i$: that is, marginal costs are zero, and inverse demand is given by $P(q) = 2 - q$.

Notice that Player i's expected payoff depends linearly on Player $-i$'s expected quantity; thus, our analysis can be restricted to degenerate (point) beliefs without loss of generality. For any $a_{-i} \in A_{-i}$, Player i's best reply is found by solving

$$\max_{a_i \in [0,1]} (2 - a_i - a_{-i})a_i, \quad \text{so} \quad r_i(a_{-i}) = 1 - \frac{1}{2}a_{-i}$$

Then notice that no $a_i \in [0, \frac{1}{2})$ is ever a best reply: after all, even if $a_{-i} = 1$, there would still be residual demand in the market, and consumers would be willing to pay 1 dollar for an additional marginal unit of the good. That is, Player i is actually facing the residual inverse demand curve $P^R(a_i) = 1 - a_i$, so setting $a_i = \frac{1}{2}$ (the corresponding monopoly quantity) is optimal.

This is true for $i = 1, 2$. Thus, assume that Player i understands that $a_{-i} \in [\frac{1}{2}, 1]$: that is, Player i believes that, since her opponent is rational, he will never produce less than $\frac{1}{2}$ units of the good. But then, producing $a_i \in (\frac{3}{4}, 1]$ units cannot be a best reply against a belief satisfying this restriction. Since, moreover, we have already concluded that quantities $a_i \in [0, \frac{1}{2})$ must be discarded on the basis of rationality considerations alone, the assumptions of

Bayesian rationality + Belief in the opponent's Bayesian rationality

already imply that only action profiles (a_1, a_2) with $a_i \in [\frac{1}{2}, \frac{3}{4}]$ should be expected.

Once we are willing to make this assumption about the players' beliefs, it becomes natural to at least conceive the possibility that *the players themselves might contemplate it.* This leads to a *hierarchy of hypotheses about mutual belief in rationality.* That is, we might think that:

(1) players are Bayesian rational;
(2) players believe that their opponents are also Bayesian rational;
(3) players believe that their opponents believe that their *own* opponents are Bayesian rational; and so on.

The following definition captures the implications of these assumptions.

Definition 4 *Fix a game $G = (N, (A_i, \mathcal{T}_i, u_i)_{i \in N})$. For every player $i \in N$, let $A_i^0 = A_i$. For $k \geq 1$ and for every $i \in I$, let $a_i \in A_i^k$ iff there exists $\alpha_{-i} \in \Delta(A_{-i}, \mathcal{B}(\mathcal{T}_{-i}))$ such that*

(1) $a_i \in r_i(\alpha_{-i})$*;*
(2) $\alpha_{-i}(\prod_{j \neq i} A_j^{k-1}) = 1$.
An action $a_i \in A_i^k$ *is said to be* k-correlated rationalizable*; an action* $a_i \in \bigcap_{k \geq 1} A_i^k$ *is said to be* correlated rationalizable.

Part (2) in the above definition captures the idea that, at each step k, we are willing to impose the assumption that players are able to perform $k-1$ steps of eductive reasoning about their opponents. Although this is not explicitly required in the definition, this corresponds to progressively more stringent assumptions about rationality; if action sets are finite, this implies that correlated rationalizable strategies exist (why?).

Proposition 0.2 *Fix a game* $G = (N, (A_i, \mathcal{T}_i, u_i)_{i \in N})$*. Then, for every* $i \in N$ *and* $k \geq 1$, $A_i^k \subset A_i^{k-1}$.

Proof: The claim is true for $k = 1$. Assume it is true for some k, fix $i \in N$ and pick $a_i \in A_i^{k+1}$. By definition, there exists $\alpha_{-i} \in \Delta(A_{-i}, \mathcal{B}(\mathcal{T}_{-i}))$ such that $a_i \in r_i(\alpha_{-i})$ and $\alpha_{-i}(A_{-i}^k) = 1$. Since by the induction hypothesis $A_{-i}^k \subset A_{-i}^{k-1}$, this implies $\alpha_{-i}(A_{-i}^{k-1}) = 1$; hence, $a_i \in A_i^k$, as claimed. ■

Independent rationalizability, or (in keeping with standard terminology) "rationalizability" tout-court, incorporates the additional restriction that players' beliefs are actually independent product measures on their opponents' action profiles; that players believe that their opponents' beliefs are independent; that they believe that their opponents believe that their respective opponents' beliefs are independent; and so on and so forth.

Definition 5 *Fix a game* $G = (N, (A_i, \mathcal{T}_i, u_i)_{i \in N})$*. For every player* $i \in N$*, let* $R_i^0 = A_i$*. For* $k \geq 1$ *and for every* $i \in I$*, let* $a_i \in R_i^k$ *iff there exists* $\alpha_{-i} \in \Delta(A_{-i}, \mathcal{B}(\mathcal{T}_{-i}))$ *such that*
(1) $a_i \in r_i(\alpha_{-i})$*;*
(2) $\alpha_{-i}(\prod_{j \neq i} R_j^{k-1}) = 1$*;*
(3) α_{-i} *is a product measure on* $A_{-i}, \mathcal{B}(\mathcal{T}_{-i})$.
An action $a_i \in R_i^k$ *is said to be* k-rationalizable*; an action* $a_i \in \bigcap_{k \geq 1} R_i^k$ *is said to be* rationalizable.

I conclude this section with two notes pertaining to the interpretation of (correlated) rationalizability.

First, I emphasize that, although we have *motivated* (correlated) rationalizability with informal assumptions about rationality and beliefs, we have been quite vague in defining what we mean by statements such as, "Player i believes that Player $-i$ is rational." We have claimed that this "implies" that Player i's beliefs should assign positive probability only to those strategies of Player $-i$ which can themselves be justified by some belief.

However, this is an informal statement. Recent contributions have shown how to construct formal models (based on the classical decision-theoretic setting) which allow one to *formalize* ideas such as "rationality" and "belief". In the framework of those models, the "implication" we have used to define rationalizability can actually be *proved*—i.e. it is a theorem. This places rationalizability on very solid methodological ground.

Second, it might seem rather natural to assume that beliefs should be independent (and that moreover this ought to be "common belief"). After all, the basic model of simultaneous games postulates that players choose their actions independently, without any possibility of coordination. Thus, correlated beliefs might even appear to be unwarranted.

The debate on "causal vs. epistemic independence," as the issue is often referred to, is lengthy and beyond the scope of these notes. Let me just invite you to reflect on the subject, based on the following example.

Consider the following three-player "coordination" game: Players 1 and 2 choose L or R, and get 1 if they coordinate and 0 if they do not. Player 3 chooses C or N; C yields 1 if Players 1 and 2 coordinate, and 0 otherwise, while N yields 0 if they coordinate and 1 if they do not. That is, Player 3's choice represents a bet on whether or not Players 1 and 2 manage to coordinate.

Now suppose that Player 3 is fully aware that Players 1 and 2 choose independently; however, she expects them to coordinate, but is unsure as to which action they will choose (i.e. L or R). Thus, she assigns probability $\frac{1}{2}$ to each of the profiles (L,L) and (R,R). This is a correlated belief, but it could be explained with the following simple story: Players 1 and 2 have been playing this game for a long time, and have "learned to coordinate." Player 3 knows this, but has never observed their actual play—perhaps because she has also been playing the game for a long time, but she has only observed her own payoff in each play of the game (and hence ascertained whether or not there was coordination).

I welcome your comments on this issue!

Rationalizability and Dominance

Consider the following definition:

Definition 6 *Fix a finite game $G = (N, (A_i, u_i)_{i \in N})$ and a player $i \in N$. An action $a_i \in A_i$ is* strictly dominated *for Player i iff there exists $\alpha_i \in \Delta(A_i)$ such that*

$$\forall a_{-i} \in A_{-i}, \quad \sum_{a_i' \in A_i} u_i(a_i', a_{-i}) \alpha_i(a_i) > u_i(a_i, a_{-i});$$

a_i *is* weakly dominated *for Player* i *iff there exists* $\alpha_i \in \Delta(A_i)$ *such that*

$$\forall a_{-i} \in A_{-i}, \quad \sum_{a_i' \in A_i} u_i(a_i', a_{-i})\alpha_i(a_i) \geq u_i(a_i, a_{-i})$$

and there exists $a_{-i} \in A_{-i}$ *such that*

$$\sum_{a_i' \in A_i} u_i(a_i', a_{-i})\alpha_i(a_i) > u_i(a_i, a_{-i}).$$

It is (relatively) easy to see that a strictly dominated strategy cannot be a best reply against any belief. The converse is also true:

Proposition 0.3 *Fix a finite game* $G = (N, (A_i, u_i)_{i \in N})$ *and a player* $i \in N$*. An action* $a_i \in A_i$ *is strictly dominated for Player* i *iff there is no belief* $\alpha_{-i} \in \Delta(A_{-i})$ *such that* $a_i \in r_i(\alpha_{-i})$*.*

Proof: Number the actions in A_i *except* a_i from 1 to $|A_i|$, and the action profiles in A_{-i} from 1 to $|A_{-i}|$. Define a $|A_i| \times |A_{-i}|$ matrix U by letting $U_{mn} = u_i(a_i^n, a_{-i}^m) - u_i(a_i, a_{-i}^m)$, where a_i^n is the n-th action in $A_i \setminus \{a_i\}$, and a_{-i}^m is the m-th action profile in A_{-i}.

Consider the problem:

$$\max_{x \geq 0} 0 \cdot x \quad \text{s.to} \quad U \cdot x \leq 0, 1'x = 1$$

where x is a $|A_{-i}|$-dimensional column vector. If such a vector exists, then it represents a belief for i, and $a_i \in r_i(x)$. The dual problem is

$$\min_{y \geq 0, w} w \quad \text{s.to} \quad y' \cdot U + w1' \geq 0$$

where y is a $|A_i - 1|$-dimensional column vector and w is a scalar (not sign-constrained). The min problem reads as follows: find a mixture y of actions for i, assigning zero weight to a_i, such that the expected payoff to the mixture is at least $-w$ higher than the payoff to a_i for any action profile of i's opponents. Note that w is unconstrained, and the problem seeks to minimize w (hence make $-w$ as large as possible). If the optimal value of the problem is 0, so $w = 0$, then a_i is not strictly dominated.

Observe that the max program is always bounded, while the min program is always feasible (take $y = 0$, $w = 0$). Thus, suppose a_i is a best reply, so the max program is feasible: then, by the Existence-Duality theorem, both problems must have optimal solutions, and the values must coincide. This implies that $w = 0$, so a_i is not strictly dominated. Suppose

instead that a_i is not a best reply to any belief, so the max program is infeasible; again, this implies that the min program must be unbounded, i.e. w can be made arbitrarily negative (positive values cannot be optimal, because 0 is a feasible value). But this means that a_i is strictly dominated: for any $K < 0$ there exists a feasible pair (y, w) such that $w \leq K$: but for this y, $\sum u_i(a_i^n, a_{-i})y^n - u_i(a_i, a_{-i}) > -K > 0$ for any $a_{-i} \in A_{-i}$. ■

Let me emphasize that it is crucial to look at domination by *mixtures* of strategies for the equivalence result to hold. If one is not willing to allow for actual randomization, then iterated strict dominance as defined above may make little sense, and one may wish to consider domination by *pure* strategies: Tilman Borgers, among others, has worked on this problem.

I will ask you for an alternative proof of the preceding proposition, based on the Separating Hyperplane theorem. I will also ask you to prove the following (easy) result, which is one reason why correlated rationalizability is worth mentioning (even if you believe in independent beliefs).

First, consider the following definition:

Definition 7 *Fix a finite game $G = (N, (A_i, u_i)_{i \in N})$. For every player $i \in N$, let $SD_i^0 = A_i$; next, for $k \geq 1$, and for every $i \in N$, say that $a_i \in SD_i^k$ iff a_i is not strictly dominated in the game $G^{k-1} = (N, (SD_i^{k-1}, u_i^{k-1})_{i \in N})$ (where u_i^{k-1} denotes the appropriate restriction of u_i).*

The idea is to discard strictly dominated strategies, then look at the game resulting from G by removing them, and iterate the definition.

Proposition 0.4 *Fix a finite game $G = (N, (A_i, u_i)_{i \in N})$. Then, for every player $i \in N$, and for every $k \geq 1$, $A_i^k = SD_i^k$.*

I will have more to say about weak dominance and iterated weak dominance when we discuss extensive games.